

GeoNeRF: Generalizing NeRF with Geometry Priors

Mohammad Mahdi Johari
Idiap Research Institute, EPFL
mohammad.johari@idiap.ch

Yann Lepoittevin
ams OSRAM
yann.lepoittevin@ams.com

François Fleuret
University of Geneva, EPFL
francois.fleuret@unige.ch




Figure 1. Our generalizable GeoNeRF model infers complex geometries of objects in a novel scene without per-scene optimization and synthesizes novel images of higher quality than the existing works: IBRNet [54] and MVSNeRF [6].

Abstract

We present *GeoNeRF*, a generalizable photorealistic novel view synthesis method based on neural radiance fields. Our approach consists of two main stages: a geometry reasoner and a renderer. To render a novel view, the geometry reasoner first constructs cascaded cost volumes for each nearby source view. Then, using a Transformer-based attention mechanism and the cascaded cost volumes, the renderer infers geometry and appearance, and renders detailed images via classical volume rendering techniques. This architecture, in particular, allows sophisticated occlusion reasoning, gathering information from consistent source views. Moreover, our method can easily be fine-tuned on a single scene, and renders competitive results with per-scene optimized neural rendering methods with a fraction of computational cost. Experiments show that *GeoNeRF* outperforms state-of-the-art generalizable neural rendering models on various synthetic and real datasets. Lastly, with a slight modification to the geometry reasoner, we also propose an alternative model that adapts to RGBD images. This model directly exploits the depth information often available thanks to depth sensors. The implementation code is available at <https://www.idiap.ch/paper/geonerf>.

1. Introduction

Novel view synthesis is a long-standing task in computer vision and computer graphics. Neural Radiance Fields

(NeRF) [37] made a significant impact on this research area by implicitly representing the 3D structure of the scene and rendering high-quality novel images. Our work addresses the main drawback of NeRF, which is the requirement to train from scratch for every scene separately. The per-scene optimization of NeRF is lengthy and requires densely captured images from each scene.

Approaches like pixelNeRF [61], GRF [51], MINE [29], SRF [8], IBRNet [54], MVSNeRF [6], and recently introduced NeRFormer [43] address this issue and generalize NeRF rendering technique to unseen scenes. The common motivation behind such methods is to condition the NeRF renderer with features extracted from source images from a set of nearby views. Despite the generalizability of these models to new scenes, their understanding of the scene geometry and occlusions is limited, resulting in undesired artifacts in the rendered outputs. MVSNeRF [6] constructs a low-resolution 3D cost volume inspired by MVSNet [58], which is widely used in the Multi-View Stereo research, to condition and generalize the NeRF renderer. However, it has difficulty rendering detailed images and does not deal with occlusions in a scene. In this work, we take MVSNeRF as a baseline and propose the following improvements.

- We introduce a geometry reasoner in the form of cascaded cost volumes (Section 3.1) and train it in a semi-supervised fashion (Section 3.4) to obtain fine and high-resolution priors for conditioning the renderer.
- We combine an attention-based model which deals with information coming from different source views

at any point in space, by essence permutation invariant, with an auto-encoder network which aggregates information along a ray, leveraging its strong Euclidean and ordering structure (Section 3.3).

- Thanks to the symmetry and generalizability of our geometry reasoner and renderer, we detect and exclude occluded views for each point in space and use the remaining views for processing that point (Section 3.3).

In addition, with a slight modification to the architecture, we propose an alternate model that takes RGBD images (RGB+Depth) as input and exploits the depth information to improve its perception of the geometry (Section 3.5).

Concurrent to our work, the followings also introduce a generalizable NeRF: RGBD-Net [38] builds a cost volume for the target view instead of source views, NeuralMVS [45] proposes a coarse to fine approach to increase speed, and NeuRay [32] proposes a method to deal with occlusions.

2. Related Work

Multi-View Stereo. The purpose of Multi-View Stereo (MVS) is to estimate the dense representation of a scene given multiple overlapping images. This field has been extensively studied: first, with now called traditional methods [9, 18, 27, 47] and more recently with methods relying on deep learning such as MVSNet [58], which outperformed the traditional ones. MVSNet estimates the depth from multiple views by extracting features from all images, aggregating them into a variance-based cost volume after warping each view onto the reference one, and finally, post-processing the cost volumes with a 3D-CNN. The memory needed to post-process the cost volume being the main bottleneck of [58], R-MVSNet [59] proposed regularizing the cost volume along the depth direction with gated recurrent units while slightly sacrificing accuracy. To further reduce the memory impact, [7, 19, 57] proposed cascaded architectures, where the cost volume is built at gradually finer scales, with the depth output computed in a coarse to fine manner without any compromise on the accuracy. Replacing the variance-based metric with group-wise correlation similarity is another approach to further decrease the memory usage of MVS networks [56]. We found MVS architectures suitable for inferring the geometry and occlusions in a scene and conditioning a novel image renderer.

Novel View Synthesis. Early work on synthesizing novel views from a set of reference images was done by blending reference pixels according to specific weights [11, 28]. The weights were computed according to ray-space proximity [28] or approximated geometry [4, 11]. To improve the computed geometry, some used the optical flow [5, 15]

or soft blending [42]. Others synthesized a radiance field directly on a mesh [10, 21] or on a point cloud [1, 35]. An advantage of these methods is that they can synthesize new views with a small number of references, but their performance is limited by the quality of 3D reconstruction [22, 46], and problems often arise in low-textured or reflective regions where stereo reconstruction tends to fail. Leveraging CNNs to predict volumetric representations stored in voxel grids [20, 24, 42] or Multi-Plane Images [16, 17, 50, 63] produces photo-realistic renderings. Those methods rely on discrete volumetric representations of the scenes limiting their outputs’ resolution. They also need to be trained on large datasets to store large numbers of samples resulting in extensive memory overhead.

Neural Scene Representations. Recently, using neural networks to represent the geometry and appearance of scenes has allowed querying color and opacity in continuous space and viewing directions. NeRF [37] achieves impressive results for novel view synthesis by optimizing a 5D neural radiance field for a scene. Building upon NeRF many improvements were made [3, 13, 30, 34, 40, 41, 44, 48, 49], but the network needs to be optimized for hours or days for each new scene. Later works, such as GRF [51], pixelNeRF [61] and MINE [29], try to synthesize novel views with very sparse inputs, but their generalization ability to challenging scenes with complex specularities is highly restricted. MVSNeRF [6] proposes to use a low-resolution plane-swept cost volume to generalize rendering to new scenes with as few as three images without retraining. Once the cost volume is computed, MVSNeRF uses a 3D-CNN to aggregate image features. This 3D-CNN resembles the generic view interpolation function presented in IBRNet [54] that allows rendering novel views on unseen scenes with few images. Inspired by MVSNeRF, our work first constructs cascaded cost volumes per source view and then aggregates the cost volumes of the views in an attention-based approach. The former allows capturing high-resolution details, and the latter addresses occlusions.

3. Method

We use volume rendering techniques to synthesize novel views given a set of input source views. Our proposed architecture is presented in Figure 2, and the following sections provide details of our method.

3.1. Geometry Reasoner

Given a set of V nearby views $\{I_v\}_{v=1}^V$ with size $H \times W$, our geometry reasoner constructs cascaded cost volumes for each input view individually, following the same approach in CasMVSNet [19]. First, each image goes through a Feature Pyramid Network (FPN) [31] to generate semantic 2D

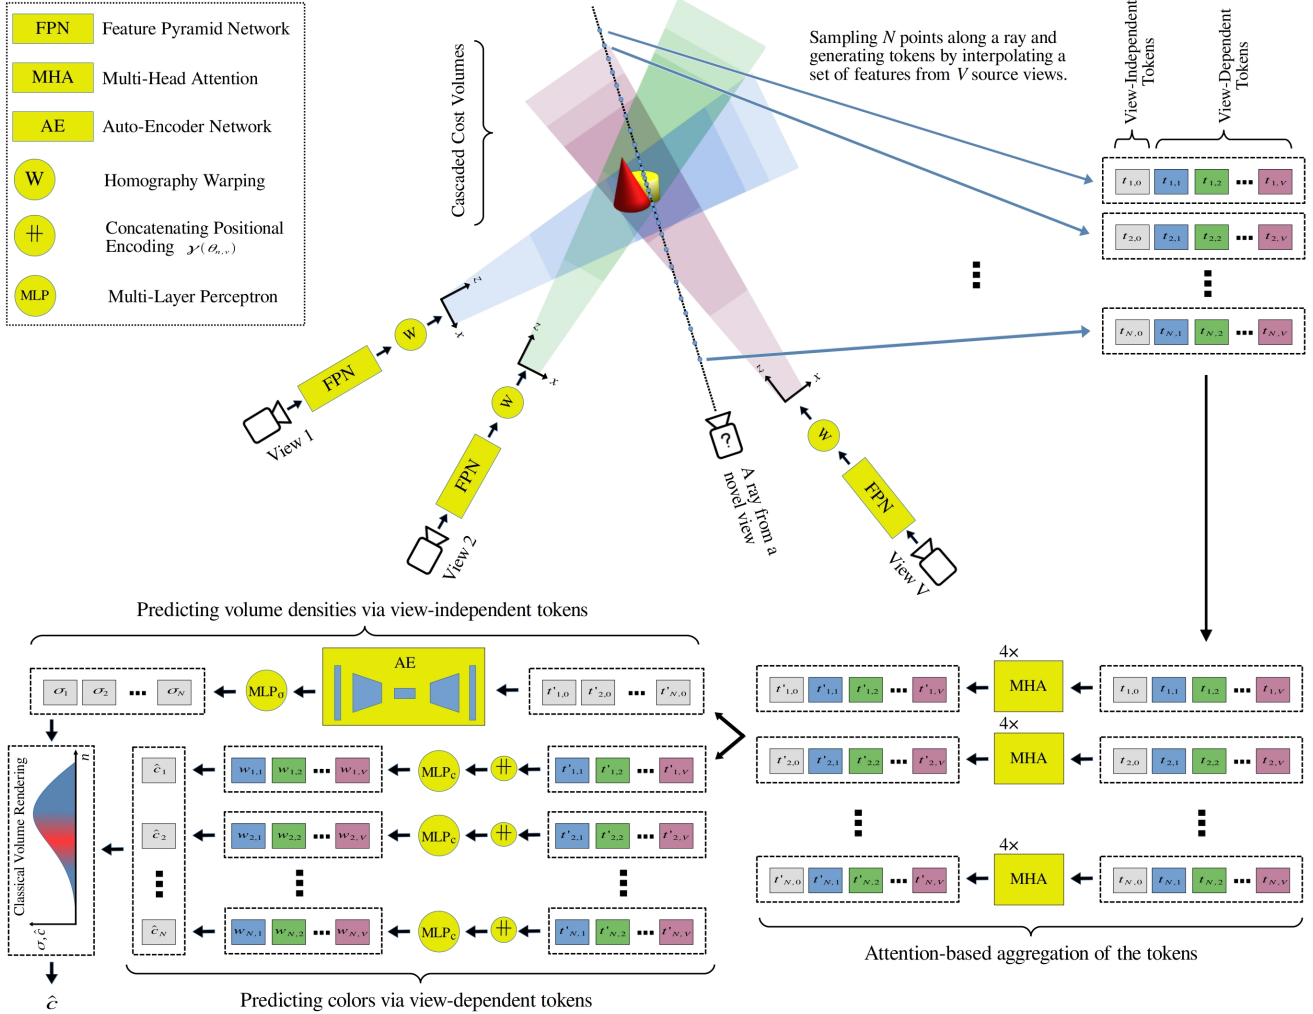


Figure 2. The overview of GeoNeRF. 2D feature pyramids are first generated via Feature Pyramid Network (FPN) [31] for each source view v . We then construct cascaded cost volumes at three levels for each view by homography warping of its nearby views (see Section 3.1). Guided by the distribution of the cascaded cost volumes in the 3D space, $N = N_c + N_f$ points $\{x_n\}_{n=1}^N$ are sampled along a ray for a novel pose (see Section 3.2). By interpolating both 2D and 3D features ($f_{n,v}^{(0)}, \{\Phi_{n,v}^{(l)}\}_{l=0}^2$) from FPN and cascaded cost volumes for each sample point x_n , one view independent token $t_{n,0}$ and V view-dependent tokens $\{t_{n,v}\}_{v=1}^V$ are generated. These $V+1$ tokens go through four stacked Multi-Head Attention (MHA) layers and yield more refined tokens $\{t'_{n,v}\}_{v=0}^V$. The MHA layers are shared among all sample points on a ray. Thereafter, the view-independent tokens $\{t'_{n,0}\}_{n=1}^N$ are regularized and aggregated along the ray samples through the AE network, and volume densities $\{\sigma_n\}_{n=1}^N$ of the sampled points are estimated. Other tokens $\{t'_{n,v}\}_{v=1}^V$, supplemented with the positional encodings $\{\gamma(\theta_{n,v})\}_{v=1}^V$, predict the color weights $\{w_{n,v}\}_{v=1}^V$ with respect to source views, and the color \hat{c}_n of each point is estimated in a weighted sum fashion (see Section 3.3). Finally, the color of the ray \hat{c} is rendered using classical volume rendering.

features at three different scale levels.

$$f_v^{(l)} = \text{FPN}(I_v) \in \mathbb{R}^{\frac{H}{2^l} \times \frac{W}{2^l} \times 2^l C} \quad \forall l \in \{0, 1, 2\} \quad (1)$$

where FPN is the Feature Pyramid Network, C is the channel dimension at level 0, and l indicates the scale level. Once 2D features are generated, we follow the same approach in CasMVSNet to construct plane sweeps and cascaded cost volumes at three levels via differentiable homography warping. CasMVSNet originally estimates depth maps $\hat{D}^{(l)}$ of the input images at three levels. The coarsest level ($l = 2$) consists of $D^{(2)}$ plane sweeps covering

the whole depth range in the camera's frustum. Then, subsequent levels narrow the hypothesis range (decrease $D^{(l)}$) but increase the spatial resolution of each voxel by creating $D^{(l)}$ finer plane sweeps on both sides of the estimated depths from the previous level. As a result, the finer the cost volume is, the thinner the depth range it covers. We make two modifications to the CasMVSNet architecture and use it as the geometry reasoner. Firstly, we provide an additional output head to the network to produce multi-level semantic 3D features $\Phi^{(l)}$ along with the estimated depth maps $\hat{D}^{(l)}$. Secondly, we replace the variance-based met-

ric in CasMVSNet with **group-wise correlation similarity** from [56] to construct light-weight volumes. Group-wise correlation decreases memory usage and inference time.

To be more specific, for each source view I_v , we first form a set of its nearby views Γ_v . Then, by constructing $D^{(l)}$ depth plane sweeps and homography warping techniques, we create multi-level cost volumes $P_v^{(l)}$ from 2D feature pyramids $f^{(l)}$ of images in Γ_v using the group-wise correlation similarity metric from [56]. We finally further process and regularize the cost volumes using 3D hourglass networks $R_{3D}^{(l)}$ and generate depth maps $\hat{D}_v^{(l)} \in \mathbb{R}^{\frac{H}{2^l} \times \frac{W}{2^l} \times 1}$ and 3D feature maps $\Phi_v^{(l)} \in \mathbb{R}^{D^{(l)} \times \frac{H}{2^l} \times \frac{W}{2^l} \times C}$:

$$\hat{D}_v^{(l)}, \Phi_v^{(l)} = R_{3D}^{(l)} \left(P_v^{(l)} \right) \quad \forall l \in \{0, 1, 2\} \quad (2)$$

3.2. Sampling Points on a Novel Ray

Once the features from the geometry reasoner are generated, we render novel views with the ray casting approach. For each camera ray at a novel camera pose, we first sample N_c points along the ray uniformly to cover the whole depth range. Furthermore, we estimate a stepwise probability density function $p_0(x)$ along the ray representing the probability that a point x is covered by a full-resolution partial cost volume $P^{(0)}$. Voxels inside the thinnest, full-resolution cost volumes $\{P_v^{(0)}\}_{v=1}^V$ contain the most valuable information about the surfaces and geometry. Therefore, we sample N_f more points from $p_0(x)$ distribution. Unlike previous works [2, 37, 43, 54, 61] that require training two networks simultaneously (one coarse and one fine network) and rendering volume densities from the coarse network to resample more points for the fine one, we sample a mixture of $N = N_c + N_f$ valuable points before rendering the ray without any computation overhead or network duplication thanks to the design of our geometry reasoner.

3.3. Renderer

For all sample points $\{x_n\}_{n=1}^N$, we interpolate the full-resolution 2D features $f_{n,v}^{(0)}$ and the three-level 3D features $\{\Phi_{n,v}^{(l)}\}_{l=0}^2$ from all source views. We also define an occlusion mask $M_{n,v}$ for each point x_n with respect to each view v . Formally, if a point x_n stands behind the estimated full-resolution depth map $\hat{D}_v^{(0)}$ (being occluded) or the projection of x_n to the camera plane of view v lies outside of the image plane (being outside of the camera frustum), we set $M_{n,v} = 0$ and discard view v from the rendering process of point x_n . Next, we create a view-independent token $t_{n,0}$ and V view-dependent tokens $\{t_{n,v}\}_{v=1}^V$ by utilizing the interpolated features for each point x_n :

$$\begin{aligned} t_{n,v} &= \text{LT} \left(\left[f_{n,v}^{(0)}; \{\Phi_{n,v}^{(l)}\}_{l=0}^2 \right] \right) \quad \forall v \in \{1, \dots, V\} \\ t_{n,0} &= \text{LT} \left(\left[\text{mean}\{f_{n,v}^{(0)}\}_{v=1}^V; \text{var}\{f_{n,v}^{(0)}\}_{v=1}^V \right] \right) \end{aligned} \quad (3)$$

where $\text{LT}(\cdot)$ and $[\cdot; \cdot]$ denote respectively linear transformation and concatenation. $t_{n,0}$ could be considered as a global understanding of the scene at point x_n , while $t_{n,v}$ represents the understanding of the scene from source view v . The global and view-dependent tokens are aggregated through four stacked Multi-Head Attention (MHA) layers, which are introduced in Transformers [14, 52]:

$$\{t'_{n,v}\}_{v=0}^V = 4 \times \text{MHA} \left(t_{n,0}, \{t_{n,v}, M_{n,v}\}_{v=1}^V \right) \quad (4)$$

Our MHA layers also take the occlusion masks $M_{n,v}$ as inputs and force the occluded views' attention scores to zero to prevent them from contributing to the aggregation.

The global view-independent output tokens $\{t'_{n,0}\}_{n=1}^N$ now have access to all necessary data to learn the geometry of the scene and estimate volume densities. We further regularize these tokens through an auto-encoder-style (AE) network in the ray dimension (n). The AE network learns the global geometry along the ray via convolutional layers and predicts more coherent volume densities σ_n :

$$\{\sigma_n\}_{n=1}^N = \text{MLP}_\sigma \left(\text{AE} \left(\{t'_{n,0}\}_{n=1}^N \right) \right) \quad (5)$$

where MLP_σ is a simple two-layer perceptron. We argue that convolutionally processing the tokens with the AE network along the ray dimension (n) is a proper inductive bias and significantly reduces the computation resources compared to methods like IBRNet [54] and NeRFormer [43], which employ an attention-based architecture because the geometry of a scene is naturally continuous, and accordingly, closer points are more likely related.

View-dependent tokens $\{t'_{n,v}\}_{v=1}^V$, together with two additional inputs, are used for color prediction. We project each point x_n to source views' image planes and interpolate the color samples $c_{n,v}$. We also calculate the angle between the novel camera ray and the line that passes through the camera center of source view v and x_n . This angle $\theta_{n,v}$ represents the similarity between the camera pose of source view v and the novel view. Each point's color is estimated via a weighted sum of the non-occluded views' colors:

$$\begin{aligned} w_{n,v} &= \text{Softmax} \left(\left\{ \text{MLP}_c \left([t'_{n,v}; \gamma(\theta_{n,v})] \right), M_{n,v} \right\}_{v=1}^V \right) \\ \hat{c}_n &= \sum_{v=1}^V w_{n,v} c_{n,v} \quad \forall n \in \{1, 2, \dots, N\} \end{aligned} \quad (6)$$

where $\gamma(\cdot)$ is the sinusoidal positional encoding proposed in [37], and MLP_c is a simple two-layer perceptron. The Softmax function also takes the occlusion masks $M_{n,v}$ as input to exclude occluded views.

Once volume densities and colors are predicted, our model renders, as in NeRF [37], the color of the camera ray at a novel pose using the volume rendering approach:

$$\hat{c} = \sum_{n=1}^N \exp \left(- \sum_{k=1}^{n-1} \sigma_k \right) (1 - \exp(-\sigma_n)) \hat{c}_n \quad (7)$$

In addition to the rendered color, our model also outputs the estimated depth \hat{d} for each ray:

$$\hat{d} = \sum_{n=1}^N \exp\left(-\sum_{k=1}^{n-1} \sigma_k\right) (1 - \exp(-\sigma_n)) z_n \quad (8)$$

where z_n is the depth of point x_n with respect to the novel pose. This auxiliary output is helpful for training and supervising our generalizable model (see Section 3.4).

3.4. Loss Functions

The primary loss function when we train our generalizable model on various scenes and the **only** loss when we fine-tune it on a specific scene is the mean squared error between the rendered colors and ground truth pixel colors:

$$\mathcal{L}_c = \frac{1}{|R|} \sum_{r \in R} \|\hat{c}(r) - c_{gt}(r)\|^2 \quad (9)$$

where R is the set of rays in each training batch and c_{gt} is the ground truth color.

DS-NeRF [12] shows that depth supervision can help NeRF train faster with fewer input views. Moreover, numerous works [25, 39, 53, 60] show that despite the high-quality color rendering, NeRF has difficulty reconstructing 3D geometry and surface normals. Accordingly, for training samples coming from datasets with ground truth depths, we also output the predicted depth \hat{d} for each ray and supervise it if the ground truth depth of that pixel is available:

$$\mathcal{L}_d = \frac{1}{|R_d|} \sum_{r \in R_d} \|\hat{d}(r) - d_{gt}(r)\|_{s1} \quad (10)$$

where R_d is the set of rays from samples with ground truth depths and d_{gt} is the pixel ground truth depth and $\|\cdot\|_{s1}$ is the smooth L_1 loss.

Lastly, we supervise cascaded depth estimation networks in our geometry reasoner. For datasets with ground truth depth, the loss is defined as:

$$\mathcal{L}_D^{(l)} = \frac{2^{-l}}{|V|} \sum_{v=1}^V \left\langle \|\hat{D}_v^{(l)} - D_v^{(l)}\|_{s1} \right\rangle \quad (11)$$

where $D_v^{(l)}$ is the ground truth depth map of view v resized to scale level l , and $\langle \cdot \rangle$ denotes averaging over all pixels. For training samples without ground truth depths, we self-supervise the depth maps. We take the rendered ray depths as pseudo-ground truth and warp their corresponding colors and estimated depths from all source views using camera transformation matrices. If the ground truth pixel color of a ray is consistent with the warped color of a source view, and it is located in a textured neighborhood, we allow \hat{d} to

supervise the geometry reasoner for that view. Formally:

$$\mathcal{L}_D^{(l)} = \frac{2^{-l}}{|V||R|} \sum_{v=1}^V \sum_{r \in R} M_v(r) \|\hat{D}_v^{(l)}(r_v) - \hat{d}(r_v)\|_{s1}$$

$$\text{where } r_v = T_{\rightarrow v}(r, \hat{d}(r)) \quad (12)$$

$$\text{and } M_v(r) = \begin{cases} 1 & \text{if } |I_v(r_v) - c_{gt}(r)| < \epsilon_c \\ & \text{and } V_5(I_v(r_v)) > \epsilon_t \\ 0 & \text{otherwise} \end{cases}$$

Given a ray r at a novel pose with rendered depth $\hat{d}(r)$, $T_{\rightarrow v}(r, \hat{d}(r))$ transforms the ray to its correspondent ray from source view v using camera matrices. $\hat{d}(r_v)$ denotes the rendered depth of the correspondent ray with respect to source view v , and $M_v(r)$ validates the texturedness and color consistency. We keep pixels whose variance $V_5(\cdot)$ in their 5×5 pixels neighborhood is higher than ϵ_t , and whose color differs less than ϵ_c from the color of the ray r . The aggregated loss function for our generalizable model is:

$$\mathcal{L} = \mathcal{L}_c + 0.1\mathcal{L}_d + \lambda \sum_{l=0}^2 \mathcal{L}_D^{(l)} \quad (13)$$

where λ is 1.0 if the supervision is with ground truth depths and is 0.1 if it is with pseudo-ground truth rendered depths. For **fine-tuning** on a single scene, regardless of the availability of depth data, we only use \mathcal{L}_c as the loss function.

3.5. Compatibility with RGBD data

Concerning the ubiquitousness of the embedded depth sensors in devices nowadays, we also propose an RGBD compatible model, GeoNeRF_{+D}, by making a small modification to the geometry reasoner. We assume an **incomplete**, low-resolution, noisy depth map $D_v \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 1}$ is available for each source view v . When we construct the coarsest cost volume $P_v^{(2)}$ with $D^{(2)}$ depth planes, we also construct a binary volume $B_v \in \mathbb{R}^{D^{(2)} \times \frac{H}{4} \times \frac{W}{4} \times 1}$ and concatenate it with $P_v^{(2)}$ before feeding them to the $R_{3D}^{(2)}$ network:

$$B_v(d, h, w) = \begin{cases} 1 & \text{if } Q(D_v(h, w)) \equiv d \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

where $Q(\cdot)$ maps and quantizes real depth values to depth plane indices. B_v plays the role of coarse guidance of the geometry in GeoNeRF_{+D}. As a result of this design, the model is robust to quality and sparsity of the depth inputs.

4. Experiments

Training datasets. We train our model on the real DTU dataset [23] and real forward-facing datasets from LLFF [36] and IBNet [54]. We exclude views with

Method	Settings	DTU MVS [†] [23]			Realistic Synthetic [37]			Real Forward Facing [36]		
		PSNR [↑]	SSIM [↑]	LPIPS [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]
pixelNeRF [61]	No per-scene Optimization	19.31	0.789	0.382	7.39	0.658	0.411	11.24	0.486	0.671
IBRNet [54]		26.04	0.917	0.190	<u>25.49</u>	<u>0.916</u>	<u>0.100</u>	<u>25.13</u>	<u>0.817</u>	<u>0.205</u>
MVSNeRF [6]		<u>26.63</u>	<u>0.931</u>	<u>0.168</u>	23.62	0.897	0.176	21.93	0.795	0.252
GeoNeRF		31.34	0.959	0.060	28.33	0.938	0.087	25.44	0.839	0.180
IBRNet [54]	Per-scene Optimization	31.35	0.956	<u>0.131</u>	28.14	0.942	0.072	26.73	0.851	0.175
MVSNeRF [‡] [6]		28.50	0.933	0.179	27.07	0.931	0.168	25.45	0.877	0.192
NeRF [37]		27.01	0.902	0.263	31.01	0.947	0.081	26.50	0.811	0.250
GeoNeRF _{10k}		31.66	0.961	0.059	<u>30.42</u>	0.956	0.055	<u>26.58</u>	<u>0.856</u>	0.162
GeoNeRF _{1k}		<u>31.52</u>	<u>0.960</u>	0.059	29.83	<u>0.952</u>	<u>0.061</u>	26.31	0.852	<u>0.164</u>

[†] The evaluations of the baselines on the DTU MVS dataset [23] are borrowed from the paper MVSNeRF [6]. Also, metrics are calculated for all methods on this dataset on foreground pixels, whose ground truth depths stand inside the scene bound.

[‡] For fine-tuning, MVSNeRF [6] discards CNNs and directly optimizes 3D features. Direct optimization without regularization severely suffers from overfitting which is not reflected in their reported accuracy because their evaluation is done on a custom test set instead of the standard one. *E.g.*, their PSNR on NeRF dataset [37] would drop from **27.07** to **20.02** if it was evaluated on the standard test set.

Table 1. Quantitative comparison of our proposed GeoNeRF with existing generalizable NeRF models in terms of PSNR (higher is better), SSIM [55] (higher is better), and LPIPS [62] (lower is better) metrics. Highlights are **best** and second best. GeoNeRF is superior to the existing approaches in all experiments in which the methods are evaluated without any per-scene optimization (the top row). Notably, GeoNeRF outperforms others with a significant margin on datasets with relatively sparser source views: DTU MVS [23] and NeRF realistic synthetic [37]. In particular, our method generalizes outstandingly well on the NeRF synthetic dataset [37], although our training dataset only contains real scenes which greatly differ from the synthetic scenes. The bottom row of the table presents the evaluation of the methods when they are fine-tuned on each scene separately, as well as a comparison with vanilla NeRF [37], which is per-scene optimized for 200k–500k iterations. After fine-tuning for only 10k iterations, our GeoNeRF_{10k} produces competitive results with NeRF. Remarkably, even after fine-tuning for 1k iterations (approximately one hour on a single V100 GPU), GeoNeRF_{1k} reaches 98.15% of the GeoNeRF_{10k}’s performance on average, which is another evidence for efficient convergence of our model on novel scenes.

incorrect exposure from the DTU dataset [23] following the practice in pixelNeRF [61] and use the same 88 scenes for training as in pixelNeRF [61] and MVSNeRF [6]. Ground truth depths of DTU [23], provided by [58], are the only data that is directly used for depth supervision. For samples from forward-facing datasets (35 scenes from LLFF [36] and 67 scenes from IBRNet [54]), depth supervision is in the self-supervised form.

Evaluation datasets. We evaluate our model on the 16 test scenes of DTU MVS [23], the 8 test scenes of real forward-facing dataset from LLFF [36], and the 8 scenes in NeRF realistic synthetic dataset [37]. We follow the same evaluation protocol in NeRF [37] for the synthetic dataset [37] and LLFF dataset [36] and the same protocol in MVSNeRF [6] for the DTU dataset [23]. Specifically, for LLFF [36], we hold out $\frac{1}{8}$ of the views of the unseen scenes, and for DTU [23], we hold out 4 views of the unseen scenes for testing and leave the rest for fine-tuning.

Implementation details. We train the generalizable GeoNeRF for 250k iterations. For each iteration, one scene is randomly sampled, and 512 rays are randomly selected as the training batch. For training on a specific scene, we only fine-tune the model for 10k iterations (GeoNeRF_{10k}), in contrast to NeRF [37] that requires 200k–500k optimiza-

tion steps per scene. Since our renderer’s architecture is agnostic to the number of source views, we flexibly employ a different number of source views V for training and evaluation to reduce memory usage. We use $V = 6$ source views for training the generalizable model and $V = 9$ for evaluation. For fine-tuning, we select V based on the images’ resolution and available GPU memory. Specifically, we set $V = 9$ for the DTU dataset [23] and $V = 7$ for the other two datasets. We fix the number of sample points on a ray to $N_c = 96$ and $N_f = 32$ for all scenes. We utilize Adam [26] as the optimizer with a learning rate of 5×10^{-4} for training the generalizable model and a learning rate of 2×10^{-4} for fine-tuning. A cosine learning rate scheduler [33] without restart is also applied to the optimizer. For the details of our networks’ architectures, refer to the supplementary.

4.1. Experimental Results

We evaluate our model and provide a comparison with the original vanilla NeRF [37] and the existing generalizable NeRF models: pixelNeRF [61], IBRNet [54], and MVSNeRF [6]. The authors of NeRFFormer [43] did not publish their code, did not benchmark their method on NeRF benchmarking datasets, nor test state-of-the-art generalizable NeRF models on their own dataset. They perform on par with NeRF in their experiments with scene-specific optimization and train and test their generalizable model on

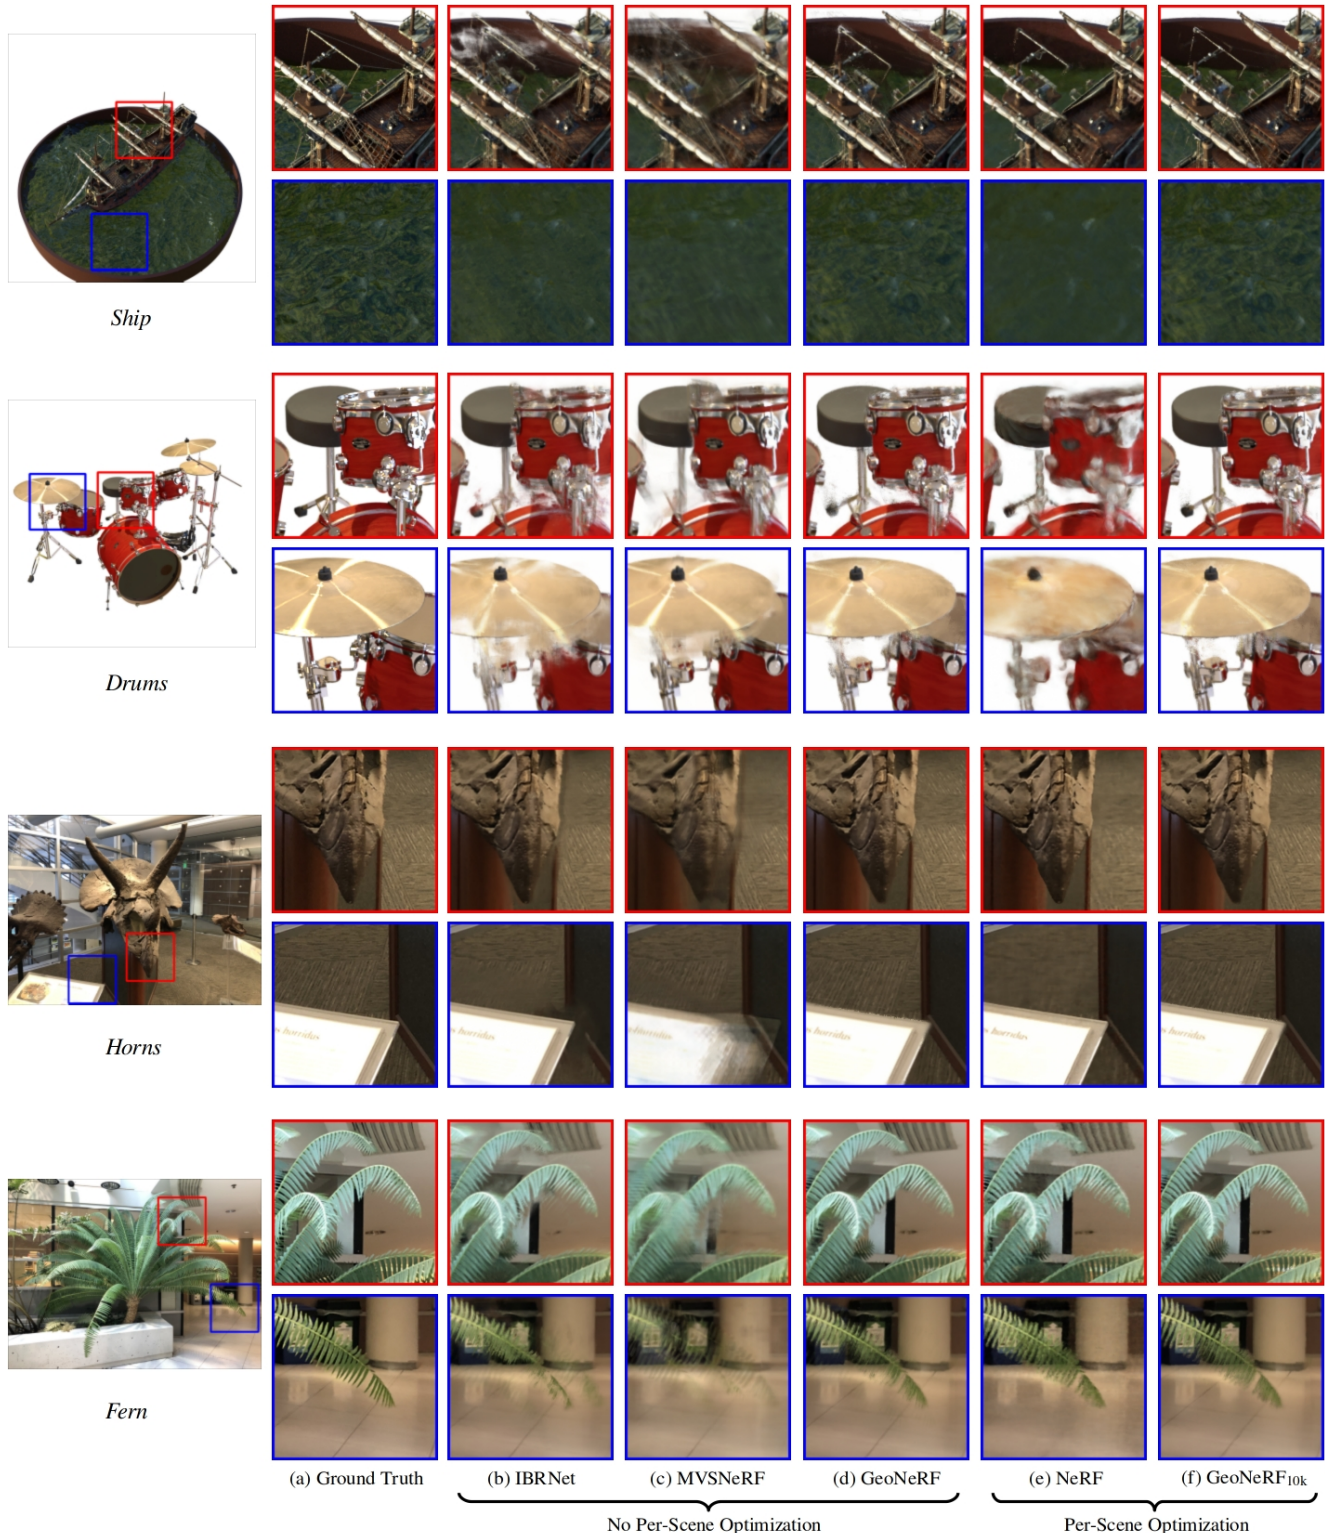


Figure 3. Qualitative comparison of the methods on the NeRF synthetic dataset [34] (*Ship* and *Drums*) and the real forward-facing dataset [36] (*Horns* and *Fern*). Our proposed GeoNeRF more accurately preserves the details of the scenes while it generates fewer artifacts than IBRNet [54] and MVSNeRF [6] (e.g. the leaves in *Fern* or the cymbal in *Drums*). After fine-tuning our model only for 10k iterations on each individual scene (GeoNeRF_{10k}), the results are competitive with per-scene optimized vanilla NeRF [37]. Compared with NeRF, GeoNeRF models produce smoother surfaces in *Drums* and higher quality textures for the water in *Ship* and for the floor in *Horns*.




Figure 4. Qualitative comparison of our generalizable GeoNeRF model with MVSNeRF [6], the state-of-the-art model on the DTU dataset [23]. Images are from DTU test scenes. Our method renders sharper images with fewer artifacts.

specific object categories separately. A quantitative comparison is provided in Table 1 in terms of PSNR, SSIM [55], and LPIPS [62]. The results show the superiority of our GeoNeRF model with respect to the previous generalizable models. Moreover, when fine-tuned on the scenes for only 10k iterations, GeoNeRF_{10k} produces competitive results with NeRF, which requires lengthy per-scene optimization. We further show that even after 1k iterations (approximately one hour on a single V100 GPU), GeoNeRF_{1k}'s results are comparable with NeRF's.

The qualitative comparisons of our model with existing methods on different datasets are provided in Figures 3 and 4. The images produced by our GeoNeRF model better preserve the details of the scene and contain fewer artifacts. For further qualitative analysis, an extensive ablation study, and limitations of our model, refer to the supplementary.

4.2. Sensitivity to Source Views

We conducted two experiments to investigate the robustness of our model to the number and quality of input source views. We first evaluated the impact of the number of source views on our model in Table 2. The results demonstrate the robustness to the sparsity of source views and suggest that GeoNeRF produces high-quality images even with a lower number of source images. Furthermore, we show that our method can operate with both close and distant source views. Table 3 shows the performance when we discard K nearest views to the novel pose and use the remaining source views for rendering. While distant source views are naturally less informative and degrade the quality, our model does not incur a significant decrease in performance.

Number of source views	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3	24.33	0.794	0.212
4	25.05	0.823	0.183
5	25.25	0.832	0.178
6	25.37	0.837	0.176
9	25.44	0.839	0.180
IBRNet [54] (10 views)	25.13	0.817	0.205

Table 2. Quantitative analysis of the robustness of our GeoNeRF to the number of input source views on the LLFF [36] test scenes, besides a comparison with IBRNet [54], which uses 10 source views.

K :	0	2	4	6	8
PSNR \uparrow	25.44	24.18	23.35	22.74	22.06
SSIM \uparrow	0.839	0.813	0.791	0.770	0.747
LPIPS \downarrow	0.180	0.212	0.235	0.253	0.276

Table 3. Quantitative analysis of the sensitivity of our GeoNeRF to discarded first K nearest neighbors on the LLFF [36] test scenes.

Model	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
GeoNeRF	31.34	0.959	0.060
GeoNeRF _{+D}	31.58	0.961	0.057

Table 4. A comparison of the performance of our RGBD compatible GeoNeRF_{+D} and original GeoNeRF on DTU [23] test scenes. For the details of this experiment, see Section 4.3.

4.3. Results with RGBD Images

To evaluate our RGBD compatible model, GeoNeRF_{+D}, we use the DTU dataset [23] to mimic the real scenario where incomplete, low-resolution depth images accompany RGB images. We feed our model the DTU images with a resolution of 600×800 , while we resize their incomplete depths to 150×200 . The comparison of the performance of GeoNeRF, with and without depth inputs is presented in Table 4. The results confirm that our GeoNeRF_{+D} model adapts to RGBD images and renders higher quality outputs.

5. Conclusion

We proposed GeoNeRF, a generalizable learning-based novel view synthesis method that renders state-of-the-art quality images for complex scenes without per-scene optimization. Our method leverages the recent architectures in the multi-view stereo field to understand the scene's geometry and occlusions by constructing cascaded cost volumes for source views. The data coming from the source views are then aggregated through an attention-based network, and images for novel poses are synthesized conditioned on these data. An advanced algorithm to select a proper set of nearby views or an adaptive approximation of the optimal number of required cost volumes for a scene could be promising extensions to our method.

Acknowledgement

This research was supported by ams OSRAM.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 696–712. Springer, 2020. 2
- [2] Relja Arandjelović and Andrew Zisserman. Nerf in detail: Learning to sample for view synthesis. *arXiv preprint arXiv:2106.05264*, 2021. 4
- [3] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5855–5864, October 2021. 2
- [4] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001. 2
- [5] Dan Casas, Christian Richardt, John Collomosse, Christian Theobalt, and Adrian Hilton. 4d model flow: Precomputed appearance alignment for real-time 4d video interpolation. In *Computer Graphics Forum*, volume 34, pages 173–182. Wiley Online Library, 2015. 2
- [6] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14124–14133, October 2021. 1, 2, 6, 7, 8
- [7] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep stereo using adaptive thin volume representation with uncertainty awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2524–2534, 2020. 2
- [8] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7911–7920, 2021. 1
- [9] Jeremy S De Bonet and Paul Viola. Poxels: Probabilistic voxelized volume reconstruction. In *Proceedings of International Conference on Computer Vision (ICCV)*, pages 418–425, 1999. 2
- [10] Paul Debevec, Yizhou Yu, and George Borshukov. Efficient view-dependent image-based rendering with projective texture-mapping. In *Eurographics Workshop on Rendering Techniques*, pages 105–116. Springer, 1998. 2
- [11] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 11–20, 1996. 2
- [12] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021. 5
- [13] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14313, October 2021. 2
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 4
- [15] Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. Montage4d: interactive seamless fusion of multiview video textures. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 1–11, 2018. 2
- [16] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2367–2376, 2019. 2
- [17] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5515–5524, 2016. 2
- [18] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009. 2
- [19] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020. 2
- [20] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Learning a neural 3d texture space from 2d exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8356–8364, 2020. 2
- [21] Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu Jiang, Leonidas J Guibas, Matthias Nießner, Thomas Funkhouser, et al. Adversarial texture optimization from rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1559–1568, 2020. 2
- [22] Michal Jancosek and Tomáš Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *CVPR 2011*, pages 3121–3128. IEEE, 2011. 2
- [23] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. 5, 6, 8
- [24] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016. 2

- [25] Berk Kaya, Suryansh Kumar, Francesco Sarno, Vittorio Ferrari, and Luc Van Gool. Neural radiance fields approach to deep multi-view photometric stereo. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1965–1977, 2022. 5
- [26] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014. 6
- [27] Vladimir Kolmogorov and Ramin Zabih. Multi-camera scene reconstruction via graph cuts. In *European conference on computer vision*, pages 82–96. Springer, 2002. 2
- [28] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, 1996. 2
- [29] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. Mine: Towards continuous depth mpi with nerf for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12578–12588, 2021. 1, 2
- [30] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. 2
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 2, 3
- [32] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, XiaoWei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. *arXiv preprint arXiv:2107.13421*, 2021. 2
- [33] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*, 2017. 6
- [34] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 2, 7
- [35] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rendering in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6878–6887, 2019. 2
- [36] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 5, 6, 7, 8
- [37] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. 1, 2, 4, 6, 7
- [38] Phong Nguyen, Animesh Karnewar, Lam Huynh, Esa Rahtu, Jiri Matas, and Janne Heikkila. Rgb-d-net: Predicting color and depth images for novel views synthesis. In *2021 International Conference on 3D Vision (3DV)*, pages 1095–1105. IEEE, 2021. 2
- [39] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021. 5
- [40] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 2
- [41] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and XiaoWei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. 2
- [42] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017. 2
- [43] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021. 1, 4, 6
- [44] Chris Rockwell, David F Fouhey, and Justin Johnson. Pixel-synth: Generating a 3d-consistent experience from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14104–14113, 2021. 2
- [45] Radu Alexandru Rosu and Sven Behnke. Neuralmvs: Bridging multi-view stereo and novel view synthesis. *arXiv preprint arXiv:2108.03880*, 2021. 2
- [46] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 2
- [47] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016. 2
- [48] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [49] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7495–7504, 2021. 2
- [50] Pratul P Srinivasan, Richard Tucker, Jonathan T Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 175–184, 2019. 2
- [51] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021. 1, 2
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 4
- [53] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 5
- [54] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 1, 2, 4, 5, 6, 7, 8
- [55] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6, 8
- [56] Qingshan Xu and Wenbing Tao. Learning inverse depth regression for multi-view stereo with correlation cost volume. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34 (07), pages 12508–12515, 2020. 2, 4
- [57] Jiayu Yang, Wei Mao, Jose M Alvarez, and Miaomiao Liu. Cost volume pyramid based depth inference for multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4877–4886, 2020. 2
- [58] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018. 1, 2, 6
- [59] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5525–5534, 2019. 2
- [60] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34, 2021. 5
- [61] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 2, 4, 6
- [62] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6, 8
- [63] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018. 2